

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
7 November 2002 (07.11.2002)

PCT

(10) International Publication Number
WO 02/089418 A1

(51) International Patent Classification ⁷ :	H04L 12/28	10/066,014	26 October 2001 (26.10.2001)	US
(21) International Application Number:	PCT/US02/12451	10/039,505	26 October 2001 (26.10.2001)	US
(22) International Filing Date:	19 April 2002 (19.04.2002)	10/046,333	26 October 2001 (26.10.2001)	US
		10/039,404	26 October 2001 (26.10.2001)	US
(25) Filing Language:	English	10/046,572	26 October 2001 (26.10.2001)	US
		10/061,564	26 October 2001 (26.10.2001)	US
(26) Publication Language:	English	10/046,640	26 October 2001 (26.10.2001)	US
		10/046,334	26 October 2001 (26.10.2001)	US
(30) Priority Data:		10/039,703	26 October 2001 (26.10.2001)	US
	60/287,068	10/039,877	26 October 2001 (26.10.2001)	US
	60/287,121	10/239,814	26 October 2001 (26.10.2001)	US
	60/287,069	10/066,217	26 October 2001 (26.10.2001)	US
	60/287,120	10/039,784	26 October 2001 (26.10.2001)	US
	60/286,918	10/066,159	26 October 2001 (26.10.2001)	US
	60/286,922	10/062,245	26 October 2001 (26.10.2001)	US
	60/287,081	10/044,182	26 October 2001 (26.10.2001)	US
	60/287,075	10/044,164	26 October 2001 (26.10.2001)	US
	60/314,088	10/068,329	26 October 2001 (26.10.2001)	US
	60/314,158			
	60/314,158			
	60/314,287			
	10/062,199			


(71) Applicant: THE BOEING COMPANY [US/US]; Seattle, WA 98124-2207 (US).

(72) Inventors: FOSTER, Michael, S.; 2841 S.W. 342nd Street, Federal Way, WA 98023 (US). DORSETT,

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR NETWORK MANAGEMENT

WO 02/089418 A1



(57) Abstract: A method and system for managing an interconnect fabric (110) that connects nodes. A network manager (115) manages an interconnect fabric or network of routing devices (e.g., interconnect fabric modules, switches, or routers) to allow source nodes (105) to transmit data to destination nodes. The network manager (115) receives registration requests from source nodes to send data to destination nodes, configures the routing devices of the network to establish a path from each source node to its destination node, and provides a virtual address to each source node. The virtual address identifies a path from the source node to the destination node. The source node sends the data to its destination node by providing the data along with the virtual address to a routing device of the network. Upon receiving the data and the virtual address, a source-side port of each routing device in the path uses the virtual address to identify a destination-side port through which the data and the virtual address are to be transmitted.



Michael, A.; 3241 S.W. 326th Street, Federal Way, WA 98023 (US). BRAATZ, James, C.; P.O. Box 3707, MC 13-08, Seattle, WA 98124-2207 (US). HUGHES, Rodney, A.; P.O. Box 3707, MC 13-08, Seattle, WA 98124-2207 (US). DAO, Turan, A.; P.O. Box 3707, MC 13-08, Seattle, WA 98124-2207 (US).

(74) Agents: ARNETT, Stephen, E. et al.; Perkins Coie LLP, P.O. Box 1247, Seattle, WA 98111-1247 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

METHOD AND SYSTEM FOR NETWORK MANAGEMENT

TECHNICAL FIELD

The described technology relates to a network manager for routing devices of an interconnect fabric.

BACKGROUND

The Internet has emerged as a critical commerce and communications platform for businesses and consumers worldwide. The dramatic growth in the number of Internet users, coupled with the increased availability of powerful new tools and equipment that enable the development, processing, and distribution of data across the Internet have led to a proliferation of Internet-based applications. These applications include e-commerce, e-mail, electronic file transfers, and online interactive applications. As the number of users of, and uses for, the Internet increases so does the complexity and volume of Internet traffic. According to UUNet, Internet traffic doubles every 100 days. Because of this traffic and its business potential, a growing number of companies are building businesses around the Internet and developing mission-critical business applications to be provided by the Internet.

Existing enterprise data networks ("EDNs") that support e-commerce applications providing services to customers are straining under the demand to provide added performance and added services. The growing customer demands for services, along with a highly competitive market, has resulted in increasingly complex ad hoc EDNs. Affordable, high-performance EDN solutions require extensive scalability, very high availability, and ease of management. These attributes are significantly compromised or completely lost as existing solutions are grown to meet the demand.

Current architectures of EDNs typically include three sub-networks: 1) a local area network (LAN) for web and database servers, 2) a computational network for application servers, and 3) a storage area network (SAN). The processing and storage elements attached to these sub-networks may have access to a wide area network (WAN) or metropolitan area network (MAN) through a bridging device commonly known as an edge switch. Each of these sub-networks typically uses a distinct

protocol and associated set of hardware and software including network interface adapters, network switches, network operating systems, and management applications. Communication through the EDN requires bridging between the sub-networks that requires active participation of server processing resources for protocol translation and interpretation.

There are many disadvantages to the current architecture of EDNs. The disadvantages result primarily because the multi-tiered architecture is fractured and complex. First, it is very difficult to integrate the disparate systems that use different communications protocols, interfaces, and so on. Second, overall performance suffers because each sub-network is managed separately, rather than being managed with comprehensive knowledge of the complete network. Third, the cost of maintaining three disparate types of network hardware and software can be high. Fourth, it is difficult to scale an architecture that uses such disparate systems. It would be desirable to have an architecture for EDNs that would be alleviate the many disadvantages of the current fractured multi-tiered architectures.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a network diagram illustrating various nodes of an example Fibre Channel fabric-based interconnect network that are inter-communicating using virtual identifiers.

Figure 2 is a flow diagram illustrating the discovery processing of a component of the interconnect fabric module in one embodiment.

Figure 3 is a flow diagram illustrating the discovery processing of the network manager in one embodiment.

Figure 4 is a flow diagram illustrating the process of establishing a path by the network manager in one embodiment.

Figure 5 is a flow diagram illustrating the processing of an identify virtual address component of the network manager in one embodiment.

Figure 6 is a flow diagram illustrating the processing of an initialize label table component of the network manager in one embodiment.

Figure 7 is a block diagram illustrating a distributed network manager in one embodiment.

Figure 8 is a flow diagram illustrating the processing of a component of an interconnect fabric module that processes reserved addresses in one embodiment.

DETAILED DESCRIPTION

A method and system for managing an interconnect fabric that connects nodes is provided. In one embodiment, a network manager manages an interconnect fabric or network of routing devices (e.g., interconnect fabric modules, switches, or routers) to allow source nodes to transmit data to destination nodes. The network manager receives registration requests from source nodes to send data to destination nodes, configures the routing devices of the network to establish a path from each source node to its destination node, and provides a virtual address to each source node. The virtual address identifies a path from the source node to the destination node. The source node sends the data to its destination node by providing the data along with the virtual address to a routing device of the network. Upon receiving the data and the virtual address, a source-side port of each routing device in the path uses the virtual address to identify a destination-side port through which the data and the virtual address are to be transmitted. The network manager configures the routing devices by setting the mappings from a source-side port to a destination-side port for each routing device in the path. The routing devices receive data via source-side ports and transmits data via destination-side ports.

In one embodiment, the network manager may be centralized or distributed. A centralized network manager may reside at one node connected to the interconnect fabric. The centralized network manager may provide configuration information to the routing devices using in-band communications or out-of-band communications. In-band communications refers to the use of the communications links connecting the ports of the routing devices. Out-of-band communications refers to the use of communications links used specifically to connect the routing devices to the network manager. A centralized network manager may alternatively reside within a routing device. Each routing device may have the capabilities to be the network manager. Upon initialization, the routing devices may coordinate to select which of the routing devices is to function as the network manager. A distributed network manager, in contrast, may have its functions performed at various manager devices connected directly to the routing devices. The network manager at each manager device can

control the routing device(s) to which it is directly connected. In addition, the network manager at each manager device can communicate with the network managers at other manager devices via in-band or out-of-band communications to coordinate control of the routing devices. In one embodiment, the distributed network manager can have different functions performed at various manager devices.

In one embodiment, the network manager identifies paths through the network from source nodes to destination nodes. The paths may be identified initially when the network manager starts up, may be identified when the network topology (e.g., the routing devices of the network and their interconnections) changes (e.g., as a result of a failure), or may be identified dynamically when a registration request is received from a source node. One skilled in the art will appreciate that various combinations of these techniques for identifying a path may be used. For example, the network manager may identify paths dynamically at registration, but may re-identify paths when the topology of the network changes. Regardless of which of these techniques are used, the network manager would typically need to know the topology of the network to identify the paths.

In one embodiment, the network manager dynamically discovers the topology of the network at initialization. The network manager can discover the topology in several different ways. The network manager can be provided with configuration information that identifies the routing devices of the network. The network manager can use this configuration information to send a message to each routing device asking which of its ports are connected to another device. The network manager can then send a query message via each connected port asking the connected-to device to identify itself and its port. From the responses to the query messages, the network manager can identify the connections (i.e., communications links) between the routing devices and thus the topology of the network. Alternatively, rather than sending a query message to each connected-to port, the routing devices upon initialization can request the connected-to devices to provide their identifications. The routing devices can then provide the identifications of the connected-to ports to the network manager. The configuration information along with the identifications of the connected-to ports describes the network topology.

In another embodiment, the network manager can dynamically discover the identifications of the routing devices by sending query messages through the ports of

the routing device to which it is directly connected. The network manager then becomes aware of each routing device that responds to the query. The network manager then sends a query message through the ports of each responding routing device. Alternatively, the network manager can send one query message to the routing device to which it is directly connected and that routing device can forward the query message via each of its ports to the routing device to which it is directly connected. Each port upon receiving the query message may send a message to the network manager with its identification along with the identification of the port to which it is directly connected.

In one embodiment, each routing device may dynamically discover which of its ports are connected to other devices (e.g., nodes or other routing devices) at initialization. Each port of a routing device may sense a characteristic of its communications link (e.g., voltage on a receive link) or may transmit a request and receive (or not receive) a response via its communications link to identify whether a device is connected. The network manager may poll each routing device for an indication of which ports of the routing device are connected to other devices. The network manager can then send a query message to each connected-to port to identify the port to which it is connected.

In one embodiment, the network manager establishes paths through the network of routing devices by configuring the ports of the routing devices along the path. The network manager may identify a path from a source node to a destination node using conventional path identification techniques. For example, the network manager may use a shortest path algorithm to identify the path with the smallest number of communications links or may use a congestion-based algorithm that factors in actual or anticipated network traffic to identify the path. The network manager then identifies a virtual address (i.e., a destination virtual address) for the identified path. The virtual address is sent by the source node along with the data to be transmitted to the destination node. The data and virtual address may be stored in a frame (e.g., Fibre Channel or InfiniBand) that has a header and a payload. The header may contain the virtual address and the payload may contain the data. The network manager then configures each source-side port of each routing device along the path to forward frames sent to the identified virtual address to the destination-side port of the routing device that is connected to the next communications link in the path. The configuration

information may be stored in a label table (described below) for the port that maps virtual addresses to destination-side ports. When a source-side port receives a frame with the identified virtual address, it then forwards the frame through the destination-side port in accordance with the configuration information.

In one embodiment, the network manager identifies a virtual address that is not currently in use by any source-side port along the path. Thus, when a source-side port receives a frame addressed with the identified virtual address, there is no ambiguity as which port of the routing device is the destination-side port. It is possible, however, that paths from two different source nodes to the same destination node may have a common sub-path. For example, the path from one source node may be through communications links A, X, Y, and Z, and the path from the other source node may be through communications links B, X, Y, and Z. In such a case, the network manager may use the same virtual address for both paths and share the terminal portion of the already-configured paths.

In one embodiment, the network manager may also establish a path between the destination node and the source node. The network manager may identify a new path or may use the same path that was identified between the source node and the destination node (but in the opposite direction). The network manager then identifies a virtual address (i.e., source virtual address) and configures the ports along the path in a manner that is analogous to the configuration of the path from the source node to the destination node. Whenever a source node sends a frame, it may include the source virtual address in the frame. When the destination node receives the frame, it can respond to the source node by sending a frame addressed to the source virtual address.

In one embodiment, the network manager may need to identify and configure a new path between a source node and a destination node. For example, the network manager may determine that, because of congestion, the required quality of service cannot be provided along the existing path or may detect a failure along the existing path. The network manager may be able to use the same virtual address to configure the new path. If the network manager uses each virtual address only once, then the network manager can use the same virtual address for the new path. If, however, the same virtual address is used to identify different paths, then it may be possible that the configuration of the new path may conflict with the configuration of another path that

uses the same virtual address. When the same virtual address can be used, then the network manager can change the path in a manner that is transparent to the source node. In particular, the network manager need not notify the source node of the change in the path. Also, if multiple destination nodes provide the same functionality, then the network manager may implement node load balancing by dynamically changing a path so that data will be sent to a different destination node. The use of these virtual addresses allows the changes to be made without changing the source and destination virtual addresses of the path.

In one embodiment, the network manager may reserve one or more virtual addresses for sending frames from a device (e.g., routing device or node) to the network manager. For example, such a frame may include a registration request from a source node. When the network manager is distributed, a routing device may detect when it has received a frame with a reserved virtual address and may forward the frame directly to the connected manager device for processing by the network manager. To provide flexibility, a frame directed to the network manager may include a combination of a reserved virtual address and another virtual address. When a routing device detects such a frame, it may determine whether it is configured to forward frames directed to the other virtual address using in-band communications. If so, the routing device forwards the frame through the destination-side port identified by the other virtual address. If the routing device is not configured for the other virtual address, then the routing device sends the frame to the network manager via out-of-band communications. For example, the routing device may send the frame to its directly connected manager device. In this way, the network manager can configure the network so that certain frames are forwarded to certain manager devices that provide certain functions or services of the network manager.

In one embodiment, a routing device is an interconnect fabric module ("IFM") with high-speed switching capabilities. An interconnect fabric module can be dynamically configured to interconnect its communications ports so that data can be transmitted through the interconnected ports. Multiple interconnect fabric modules can be connected to form an interconnect fabric through which nodes (e.g., computer systems) can be interconnected. In one embodiment, data is transmitted through the interconnect fabric as frames such as those defined by the Fibre Channel standard. Fibre Channel is defined in ANSI T11 FC-PH, FC-PH-2, FC-PH-3, FC-PI, and FC-FS

industry standard documents which are hereby incorporated by reference. One skilled in the art will appreciate, however, that the described techniques can be used with communications standards other than Fibre Channel. In particular, the described techniques can be used with the InfiniBand standard, which is described in the InfiniBand Architecture Specification, Vols. 1-2, Release 1.0, October 24, 2000, which is hereby incorporated by reference. The interconnect fabric module may allow the creation of an interconnect fabric that is especially well suited for interconnecting devices utilizing multiple information types such as might be required by the devices of an enterprise data network ("EDN").

In one embodiment, a virtual address may be part of a "virtual identifier" (e.g., source or destination identifier) that includes a domain address. A destination identifier of a frame may be set to a virtual identifier. The destination identifiers of the frames received by the interconnect fabric modules are used to forward the frame. Each interconnect fabric module is assigned a domain address. The interconnect fabric modules that are assigned the same domain address are in the same domain. The interconnect fabric modules use of the domain addresses to forward frames between domains. The network manager may configure the interconnect fabric modules with inter-domain paths. When an interconnect fabric module receives a frame with a destination domain address that matches its domain address, then the frame has arrived at its destination domain. The interconnect fabric module then forwards the frame in accordance with the destination virtual address since it has arrived at its destination domain. If, however, the domain addresses do not match, then the frame has not arrived at its destination domain. The interconnect fabric module forwards the frame using an inter-domain path. Each port of an interconnect fabric module may have a domain address table (configured by the network manager) that maps the domain addresses to the destination port through which frames with that domain address are to be forwarded. Thus, an interconnect fabric module may selectively use virtual addresses and domain addresses when forwarding frames.

In one embodiment, an interconnect fabric module uses a crosspoint switch to switch connect its source and destination ports. When the crosspoint switch has more switch ports than ports of the interconnect fabric module, the extra switch port can be used for administrative functions of the network manager. When an interconnect fabric module receives a frame directed to a virtual address reserved for administrative

services of the network manager, the interconnect fabric module connects the source port to the extra switch port which is connected to a manager device. When the frame is transmitted from the source port, the network manager at the manager device receives the frame and processes it in accordance with its administrative functions. In this way, administrative frames can be directly forwarded to the network manager when they are first received by an interconnect fabric module from a node.

In some embodiments, one or more virtual identifier ("VI") Network Interface Controller ("NIC") facilities on each node (e.g., one VI NIC for each network interface) facilitate the use of virtual identifiers in communicating data. When a VI NIC on a node receives an indication that a data communication to one or more remote nodes is to occur, such as from an application executing on the node, the VI NIC will identify an appropriate transmittal virtual identifier that can be used to route the data communication through the network to the appropriate remote destination nodes without being assigned to or directly associated with those destination nodes. Such data communications can include both transitory connectionless transmittals of data (e.g., unidirectional transmittals from a source to a destination) and non-transitory connections that allow multiple distinct transmittals of data (e.g., a persistent dedicated connection that allows a connection-initiating source and a connection destination to transmit data back and forth).

The VI NIC can identify an appropriate transmittal virtual identifier for routing a data communication in various ways. In some embodiments, the VI NIC will register some or all outgoing data communications with a network manager for the network, and will receive an appropriate transmittal virtual identifier to be used for that communication from the network manager. If an indicated data communication corresponds to a previously registered data communication (e.g., to an existing connection or to a previous communication to the same destination and in the same transmission manner), however, the VI NIC could instead in some embodiments use the previously received transmittal virtual identifier for that data communication rather than perform an additional registration for the indicated data communication. The manners in which a data communication can be transmitted vary with the transmission characteristics that are supported by a network, and can include factors such as a particular Class Of Service ("COS") or transmission priority.

In some embodiments, when a data communication indicated by a source can result in bi-directional communication (e.g., a response from one or more of the destinations), the VI NIC also identifies a response virtual identifier that can be used for routing data from one or more of the destinations back to the source. If the VI NIC registers the data communication with a network manager, this response virtual identifier may be received from the network manager. After identifying this response virtual identifier, the VI NIC associates it with information indicating how to process received data communications that are routed using the response virtual identifier. In some embodiments, such received data communications are processed by forwarding the data communications to one or more resources associated with the destination node, such as an executing application program, a file on storage, or a device that is part of the node. For example, if a source application on a source node initiates a bi-directional communication, a VI NIC for the source node may associate the response virtual identifier with that source application so that received responses can be forwarded to that source application (which then becomes the destination application for those received communications).

For illustrative purposes, some embodiments are described below in which the VI NIC is used as part of a Fibre Channel or InfiniBand network and/or as part of an EDN architecture. However, those skilled in the art will appreciate that the techniques of the invention can be used in a wide variety of other situations and with other types of networks, and that the invention is not limited to use in Fibre Channel or InfiniBand networks or with EDN architectures.

Figure 1 is a network diagram illustrating various nodes of an example Fibre Channel fabric-based interconnect network that are inter-communicating using virtual identifiers. In this example embodiment, multiple interconnect fabric modules ("IFMs") 110 with high-speed switching capabilities are used as intermediate routing devices to form an interconnect fabric, and multiple nodes 105, a network manager 115 and a Multi-Protocol Edge Switch ("MPEX") 120 are connected to the fabric. Each of the nodes has at least one VI NIC that uses virtual identifiers when communicating and receiving data. The MPEX is used to connect the Fibre Channel or InfiniBand network to an external network, such as an Ethernet-based network, and similarly includes at least one VI NIC. Data is transmitted through the interconnect fabric using frames such as those defined by the Fibre Channel or InfiniBand standards.

TOPOLOGY DISCOVERY

As described above, the network manager may dynamically discover the topology of the network using various different techniques. In the embodiment described below, each interconnect fabric module identifies which of its ports are connected to other devices. The network manager uses this information to send a message through each port that is connected to another device to identify the connected-to device. Figure 2 is a flow diagram illustrating the discovery processing of a component of the interconnect fabric module in one embodiment. Each port of an interconnect fabric module identifies whether it is connected to a port of another device, such as another switch or a node. The interconnect fabric module then provides to the network manager an indication of which of its ports are connected to other ports to assist in the discovery process. In blocks 201-204, the component determines whether each port is currently connected to another port. In block 201, the component selects the next port. In decision block 202, if all the ports have already been selected, then the component completes, else the component continues at block 203. In decision block 203, the component determines whether the selected port is connected to another port. This determination may be made based on various voltage levels of the communications links. If there is a connection, then the component continues at block 204, else the component loops to block 201 to select the next port of the interconnect fabric module. In block 204, the component notes the selected port as connected to another port and loops to block 201 to select the next port of the interconnect fabric module.

Figure 3 is a flow diagram illustrating the discovery processing of the network manager in one embodiment. The network manager first retrieves an indication of which ports of the interconnect fabric modules are connected to other devices. The network manager then sends a query message through each of the indicated ports to the connected-to port. When the connected-to port receives the query message, it responds with an identification of its interconnect fabric module and its port number. In this way, the network manager can discover the topology of the interconnect fabric. In blocks 301-303, the network manager retrieves the indications of which ports of the interconnect fabric modules are connected to other ports. In block 301, the network manager selects the next interconnect fabric module that has not yet been selected. In decision block 302, if all the interconnect fabric modules have already been selected,

then the network manager continues at block 304, else the network manager continues at block 303. In block 303, the network manager retrieves an indication of which ports of the selected interconnect fabric module are connected to other ports. The network manager may send the message using either in-band or out-of-band communications. The network manager then loops to block 301 to select the next interconnect fabric module. In blocks 304-310, the network manager determines the identity of each of the connected-to ports. In block 304, the network manager selects the next interconnect fabric module. In decision block 304, if all the interconnect fabric modules have already been selected, then the network manager completes its discovery process, else the network manager continues at block 306. In blocks 306-310, the network manager loops sending a query message through each port of the selected interconnect fabric module that is connected to another port. In block 306, the network manager selects the next port of the selected interconnect fabric module that is connected to another port. In decision block 307, if all such ports are already selected, then the network manager loops to block 304 to select the next interconnect fabric module, else the network manager continues at block 308. In block 308, the network manager sends a query message through the selected port of the selected interconnect fabric module. In block 309, the network manager receives the identification of the connected-to port of the selected port of the selected interconnect fabric module. The identification may include an indication of the interconnect fabric module and the port number of the connected-to port. In block 310, the network manager stores a mapping between the selected port of the selected interconnect fabric module and the connected-to port of the connected-to interconnect fabric module. These mappings define the topology of the network. The network manager then loops to block 306 to select the next port of the selected interconnect fabric module that is connected to another device.

The processing of the discovery of the network manager as described above assumes that the network manager initially is aware of all interconnect fabric modules of the interconnect fabric. One skilled in the art will appreciate that the network manager may become aware of additional interconnect fabric modules during the discovery process. For example, if the network manager is centralized, then it may initially send a query message through its port that is connected to the interconnect fabric. The receiving port responds with the identity and interconnect fabric module

and its port number. The network manager can then request that identified the interconnect fabric module to provide an indication of which of its ports are connected to other ports. The network manager can then send a query message through each of the indicated ports to the connected-to ports. The connected-to ports then respond with the identification of the connected-to interconnect fabric module and connected-to port. This process can be repeated transitively by the network manager to identify all interconnect fabric modules that comprise the interconnect fabric.

ESTABLISHING A PATH

Figure 4 is a flow diagram illustrating the process of establishing a path by the network manager in one embodiment. A path is typically established when a node registers with the network manager. An establish path component of the network manager may receive an indication of a source node and a destination and then identify paths of ports of interconnect fabric modules from the source node to the destination node and from the destination node to the source node. The component then identifies virtual addresses for the paths and initializes the label tables of the ports of the interconnect fabric modules along the identified paths. A label table of a port contains mappings from virtual addresses to destination-side ports through which a frame sent to that virtual address is to be forwarded. In block 401, the component identifies the paths. In one embodiment, the path from the source node to the destination node and the path from the destination node to the source node use the same ports of the same interconnect fabric modules. That is, the paths use the same communications links. Alternatively, the path in one direction may be different from the path in the other direction. One skilled in the art will appreciate that various well-known techniques for identifying paths can be used. In block 402, the component invokes an identify virtual address component passing an indication of the path and an indication that the virtual address to be used by the source node when sending a communications to the destination node (e.g., the destination virtual address). The invoked component may select a virtual address that is not currently in use by any of the source-side ports of the path. A source-side port of the path is a port that receives data sent by a source node, and a destination-side port of the path is a port through which data is transmitted on its way to the destination node. In block 403, the component invokes an identify virtual address component passing an indication of the

path and that the virtual address is to be used by the destination node (e.g., the source virtual address). In block 404, the component invokes a component to initialize the label tables of the source-side ports of the path with the destination virtual address. The invoked component transmits instructions to each source-side port of the path indicating that the port is to update its label table to map the source virtual address to a destination-side port of the interconnect fabric module. In block 405, the component invokes a component to initialize the label tables of the destination-side ports of the path with the source virtual address. The component then completes.

Figure 5 is a flow diagram illustrating the processing of an identify virtual address component of the network manager in one embodiment. In this embodiment, the identify virtual address component is provided an indication and a path along with an indication of whether a virtual address for the source node or the destination node is to be identified. The component may check every port along the path to identify a virtual address that is not currently used by a port along the path. Alternatively, the component may identify virtual addresses based on a sequential ordering. That is, the component may keep track of the last identified virtual address and increment that virtual address to identify the next virtual address. In this way, each virtual address is unique. In blocks 501-505, the component loops selecting the next virtual address and determining whether it is available. The virtual address may not be available to a port along the path when that port already uses that virtual address. In blocks 501, the component selects to the next virtual address. In decision block 502, if all the virtual addresses have already been selected, then the component indicates that a virtual address could not be identified, else the component continues at block 503. In blocks 503-505, the component loops selecting each port along the path and determining whether that port already uses the selected virtual address. In block 503, the component selects the next interconnect fabric module and port of the path. In decision block 504, if all the interconnect fabric modules and ports of the path have already been selected, then the component uses the selected virtual address as the identified virtual address and then completes, else the component continues at block 505. In decision block 505, if the selected virtual address is available at the selected interconnect fabric module and selected port, then the component loops to block 503 to select the next port along the path, else the component loops to block 501 to select the next virtual address.

Figure 6 is a flow diagram illustrating the processing of an initialize label table component of the network manager in one embodiment. The initialize label table component sends a command to each port along the path indicating to add a mapping from the identified virtual address to the other port of that interconnect fabric module along the path. The component is passed in indication of the path, the virtual address, and an indication of whether the virtual address is a source virtual address or a destination virtual address. In block 601, the component selects the next interconnect fabric module and port in the path based on whether the source or destination virtual address has been passed. In decision block 602, if all the interconnect fabric modules along the path have already been selected, then the component completes, else the component continues at block 603. In block 603, the component sends a message to be selected port of the interconnect fabric module indicating to add to its label table a mapping from the virtual address to the other port of the path. The component then loops to block 601 to select the next interconnect fabric module and port in the path.

RESERVED ADDRESSING

In one embodiment, the crosspoint switch of an IFM may have more outputs than the number of ports of the IFM. For example, a crosspoint switch may have 34 inputs and outputs, but the IFM may have only 32 ports. The IFM may use these additional ports of the crosspoint switch to route upper layer protocol frames, such as frames directed into a name server or other administrative services. In one embodiment, the additional output ports of the crosspoint switch may be connected to a manager device for the IFM. An interconnect fabric module may have a list of "reserved" addresses that designate an upper layer protocol port. When an IFM determines that an address of its frame matches one of the reserved addresses, it enables the routing of that frame to an upper layer protocol port. The routing to upper layer protocol ports may use the same arbitration mechanism as used for routing to non-upper layer protocol ports as described in the Patent Application entitled "Interconnect Fabric Module." Alternatively, when the crosspoint switch does not have extra output for an upper layer protocol port, an output can be selectively switched between a communications port and an upper layer protocol port depending on whether the address of the destination identifier is reserved.

Figure 7 is a block diagram illustrating a distributed network manager in one embodiment. In this embodiment, the network manager may be implemented on a series of manager devices connected directly to the interconnect fabric modules. The distributed network manager may communicate with each other using in-band communication of the interconnect fabric or using out-of-band communication that is independent of the interconnect fabric. The crosspoint switch of an interconnect fabric module may have reserved ports for the distributed network manager. When an interconnect fabric module receives data that designates one of the reserved ports, then the interconnect fabric module forwards the data to the distributed network manager through the reserved port.

Figure 8 is a flow diagram illustrating the processing of a component of an interconnect fabric module that processes reserved addresses in one embodiment. This component forwards the frame to the network manager via either in-band or out-of-band communications. With the use of in-band communications the frame can be routed to the appropriate interconnect fabric module, which can then send the frame to the network manager using the out-of-band communications. In block 801, if the virtual address of the received frame is a reserved address, then the component continues at block 802, else the component completes. In decision block 802, if the virtual address parameter within the frame is in the label table, then the frame is to be forwarded using in-band communications and the component continues at block 804, else the frame is to be forwarded directly to the network manager at the IFM's manager device using out-of-band communications and the component continues at block 803. In block 803, the component forwards frame to the administrative port and then completes. In block 804, the component forwards the frame based on the port map of the label table and then completes.

One skilled in the art will appreciate that, although various embodiments of the technology have been described, various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

1. A method for identifying topology of a network, the network including a plurality of routing devices, each routing device having ports, the method comprising:
 - retrieving an indication of which of the ports of the routing devices are connected to another port;
 - for each port that is connect to another port,
 - sending a query message through that port to the other port; and
 - receiving a response from the other port identifying the other device and the other port.
2. The method of claim 1 including generating a mapping from each routing device and port to device and port to which it is connected to indicate the topology of the network.
3. The method of claim 1 wherein a routing device is a switch.
4. The method of claim 1 wherein a routing device is an interconnect fabric module.
5. The method of claim 1 wherein the routing devices use virtual addresses to route frames.
6. The method of claim 1 wherein the identification of the topology is performed by a network manager.
7. The method of claim 6 wherein the network manager is distributed to the routing devices.
8. The method of claim 1 wherein the query message is sent via out-of-band communications.

9. The method of claim 1 wherein the query message is sent via in-band communications.

10. The method of claim 1 wherein the routing devices of the network are identified via the received responses.

11. The method of claim 10 wherein when a routing device is identified, retrieving an indication of which of the ports of the routing device are connected to another port.

12. The method of claim 1 wherein the retrieving of an indication of which of the ports of the routing devices are connected to another port includes sending a request to the routing device.

13. The method of claim 1 wherein the retrieving of an indication of which ports of the routing devices are connected to another port includes receiving a message from the routing device.

14. The method of claim 1 wherein each routing device determines which of its ports are connected to another port and the retrieving of an indication of which of the ports of the routing devices are connected to another port includes transmitting the determined information to a network manager.

15. A network manager for identifying topology of a network, the network including a plurality of routing devices, each routing device having ports, comprising:
a component that retrieves indications of which of the ports of the routing devices are connected to another port; and
a component that sends a query message through each port that is indicated as connected to another port to the other port and that receives a response from the other port identifying the other device and the other port.

16. The network manager of claim 15 including a component that generates a mapping from each routing device and port to the device and port to which it is connected to indicate the topology of the network.
17. The network manager of claim 15 wherein a routing device is a switch.
18. The network manager of claim 15 wherein a routing device is an interconnect fabric module.
19. The network manager of claim 15 wherein the routing devices use virtual addresses to route messages.
20. The network manager of claim 19 including a component that configures each routing device with routing data for virtual addresses.
21. The network manager of claim 20 wherein each frame of data identifies a destination virtual address.
22. The network manager of claim 15 wherein the query message is sent via out-of-band communications.
23. The network manager of claim 15 wherein the query message is sent via in-band communications.
24. The network manager of claim 15 wherein the routing devices of the network are identified via the received responses.
25. The network manager of claim 24 wherein the component that retrieves an indication of which of the ports of the routing device are connected to another port retrieves the indication when a routing device is identified.
26. The network manager of claim 25 wherein the component that retrieves an indication sends a request to a routing device.

27. The network manager of claim 15 wherein the component that retrieves an indication of which ports of the routing devices are connected to another port includes receiving a message from the routing device.

28. The network manager of claim 15 wherein each routing device determines which of its ports are connected to other ports and the retrieving of an indication of which of the ports of the routing devices are connected to another port includes receiving the determinations from the routing devices.

29. A method for identifying topology of a network, the network including a plurality of switches, each switch having ports, each port of a switch either being connected to another port or not connected to another port, the method comprising:

under control of each switch, determining whether each port of the switch is connected to a connected-to port; and

under control of a network manager,

for each of the switches,

retrieving an indication of which of the ports of the switch are connected to a connected-to port; and

for each port that is connect to a connected-to port,

sending a query message through that port to the connected-to port; and

receiving a response from the connected-to port identifying the connected-to device and connected-to port

wherein mappings from each switch and port to its connect-to device and connected-to port indicates the topology of the network.

30. The method of claim 29 wherein processing of the network manager is distributed to the switches.

31. The method of claim 29 wherein the query message is sent via out-of-band communications.

32. The method of claim 29 wherein the sending of the connect-to query message is sent via in-band communications of the network.

33. The method of claim 32 wherein the network manager identifies switches of the network via the received responses.

34. The method of claim 33 wherein when a switch is identified, the network manager performs the retrieving of the indications of which of the ports of the switch are connected to a connected-to port.

35. The method of claim 29 wherein the connected-to device is a node.

36. The method of claim 29 wherein the connected-to device is a switch.

37. A network manager for identifying topology of a network, the network including a plurality of routing devices, each routing device having ports, comprising:

means for retrieving indications of which of the ports of the routing devices are connected to another port; and

means for sending a query message through each port that is indicated as connected to another port to the other port and that receives a response from the other port identifying the other port.

38. The network manager of claim 37 including a component that generates a mapping from each port to its connected-to port to indicate the topology of the network.

39. The network manager of claim 37 wherein a routing device is a switch.

40. The network manager of claim 37 wherein a routing device is an interconnect fabric module.

41. The network manager of claim 37 wherein the routing devices use virtual addresses to route messages.

42. A method in a computer system for reconfiguring a path between a source node and a destination node, the method comprising:

establishing a first path between the source node and the destination node, the path having a virtual address;

providing the virtual address to the source node for use in transmitting data from the source node to the destination node via the established path; and

after providing the virtual address to the source node, establishing a second path between the source node and the destination node so that when the source node transmits data using the provided virtual address the data is transmitted via the second path rather than via the first path.

43. The method of claim 42 wherein the establishing of the second path is performed transparently to the source node.

44. The method of claim 42 wherein the path is established through a network of switches.

45. The method of claim 42 wherein the path is established through switches with ports and wherein the establishing of a path includes identifying a source-side port and a destination-side port for each switch.

46. The method of claim 45 wherein the establishing of the path includes providing the virtual address to each source-side port of a switch in the path.

47. The method of claim 46 wherein the virtual address of source-side port is used to map the source-side port to the destination-side port of the switch.

48. The method of claim 42 including identifying a virtual address for sending data from the source node to the destination node, the identified virtual address being provided to the source node.

49. The method of claim 48 wherein the identified virtual address is not currently used by any source-side ports of the switches.

50. The method of claim 48 wherein each port of each switch has a virtual address table for mapping virtual addresses to another port of the switch.

51. The method of claim 42 wherein when data is received at a port of a switch, the virtual address of the data is used to retrieve an indication of another port and the data is sent out of the switch through the other port.

52. The method of claim 42 wherein the establishing of path from the source node to the destination node includes identifying a source-side port and a destination-side port of each switch in the path.

53. The method of claim 42 wherein the data is a Fibre Channel frame.

54. The method of claim 42 wherein the switches are Fibre Channel compatible.

55. The method of claim 42 wherein the switches are interconnect fabric modules.

56. A computer system for reconfiguring a path between a source node and destination nodes, comprising:

a component that establishes a first path between the source node and a first destination node, the path having a virtual address, the first path being identified by a virtual address, so that when the source node transmits data using the virtual address, the data is transmitted via the first path; and

a component that, after establishing the first path, establishes a second path between the source node and a second destination node, the second path being identified by the virtual address so that when the source node transmits data using the provided virtual address after the second path is established, the data is transmitted via the second path.

57. The computer system of claim 56 including:

a component that provides the virtual address to a source node for use in transmitting data via the first path before the second path is established and via the second path after the second path is established.

58. The computer system of claim 56 wherein the establishing of the second path is performed transparently to the source node.

59. The computer system of claim 56 wherein the path is established through a network of switches.

60. The computer system of claim 56 wherein the paths are established through switches with ports and wherein the establishing of a path includes identifying a source-side port and a destination-side port for each switch in the path.

61. The computer system of claim 60 wherein the virtual address is used by source-side ports to map the source-side port to the destination-side port of the switch.

62. The computer system of claim 56 including:

a component that identifies a virtual address for sending data from the source node to a destination node, the identified virtual address being provided to the source node.

63. The computer system of claim 62 wherein the identified virtual address is not currently used by any source-side ports of the switches.

64. The computer system of claim 62 wherein each port of each switch has a virtual address table for mapping virtual addresses to another port of the switch.

65. The computer system of claim 56 wherein when data is received at a port of a switch, the virtual address of the data is used to retrieve an indication of another port and the data is sent out of the switch through the other port.

66. The computer system of claim 56 wherein the data is a Fibre Channel frame.

67. The computer system of claim 56 wherein the data is an InfiniBand frame.

68. The computer system of claim 56 wherein the first destination node and the second destination node are different nodes.

69. The computer system of claim 56 wherein the first destination node and the second destination node are the same node.

70. A computer system for reconfiguring a path between a source node and a destination node, comprising:

means for establishing a first path between the source node and the destination node, the path having a virtual address; and

means for establishing a second path between the source node and the destination node so that data transmitted using the virtual address is routed via the first path before the second path is established and via the second path after the second path is established.

71. The computer system of claim 70 including:

means for providing the virtual address to the source node for use in transmitting data to the destination node.

72. The computer system of claim 70 wherein the establishing of the second path is performed transparently to the source node.

73. The computer system of claim 70 wherein the path is established through a network of switches.

74. The computer system of claim 73 wherein the paths are established through switches with ports and wherein the means for establishing of a path includes identifying a source-side port and a destination-side port for each switch in the path.

75. The computer system of claim 74 wherein the virtual address is used by source-side ports to map the source-side port to the destination-side port of the switch.

76. The computer system of claim 73 wherein the switches are interconnect fabric modules.

77. The computer system of claim 70 including:

means for identifying a virtual address for sending data from the source node to the destination node and means for providing the virtual address to the source node.

78. The computer system of claim 77 wherein the identified virtual address is not currently used by any source-side ports of switches of the path.

79. The computer system of claim 77 wherein each port of each switch has a virtual address table for mapping virtual addresses to another port of the switch.

80. The computer system of claim 70 wherein the path comprises switches with ports and when data is received at a port of a switch, the virtual address of the data is used to retrieve an indication of another port and the data is sent out of the switch through the other port.

81. The computer system of claim 70 wherein the data is a Fibre Channel frame.

82. The computer system of claim 70 wherein the data is an InfiniBand frame.

83. A method in a computer system for establishing a path between a source node and a destination node, the method comprising:

identifying ports of switches forming a path between the source node and the destination node, each switch of the path having a source-side port and a destination-side port;

identifying a virtual address for sending data from the source node to the destination node such that the virtual address is not currently used by any of the source-side ports; and
setting each of the source-side ports to switch data sent to the identified virtual address through the destination-side port of its switch.

84. The method of claim 83 including:

identifying a virtual address for sending data from the destination node to the source node such that the virtual address is not currently used by any of the destination-side ports; and
setting each of the destination-side ports to switch data sent to the identified virtual address through the source-side port of its switch.

85. The method of claim 83 wherein each port of each switch has a virtual address table for mapping virtual addresses to another port of the switch.

86. The method of claim 83 wherein when data is received at a port of a switch, the virtual address of the data is used to retrieve an indication of another port and the data is sent out of the switch through the other port.

87. The method of claim 83 wherein a path is established between the source node and each of a plurality of destination nodes by identifying ports of switches for each path.

88. The method of claim 83 wherein the data is a Fibre Channel frame.

89. The method of claim 83 wherein the switches are Fibre Channel compatible.

90. The method of claim 83 wherein the switches are interconnect fabric modules.

91. The method of claim 83 wherein when a port of a switch receives data with a virtual address that has not been set for the port, the port does not forward the data.

92. A method for establishing a path between a source node and a destination node through a network of routing devices, the method comprising:

identifying ports of routing devices forming a path between the source node and the destination node, each routing device of the path having an identified source-side port and an identified destination-side port;

identifying a virtual address for sending data from the source node to the destination node; and

setting each of the identified source-side ports to route data sent to the identified virtual address through the identified destination-side port of its routing device.

93. The method of claim 92 including:

identifying a virtual address for sending data from the destination node to the source node; and

setting each of the identified destination-side ports to route data sent to the identified virtual address through the identified source-side port of its routing device.

94. The method of claim 92 wherein a routing device is a switch.

95. The method of claim 92 wherein each routing device has a virtual address table for mapping virtual addresses to another port of the routing device.

96. The method of claim 92 wherein when data is received at a port of a routing device, the virtual address of the data is used to retrieve an indication of another port and the data is sent out of the routing device through the other port.

97. The method of claim 92 wherein a path is established between the source node and each of a plurality of destination nodes by identifying ports of routing devices for each path.

98. The method of claim 92 wherein the data is a Fibre Channel frame.

99. The method of claim 92 wherein the data is an InfiniBand frame.

100. The method of claim 92 wherein the routing devices are interconnect fabric modules.

101. The method of claim 92 wherein when a routing device receives data with a virtual address that has not been set for the routing device, the routing device does not forward the data.

102. The method of claim 92 wherein the identified virtual address is not currently used by any of the identified source-side ports.

103. The method of claim 92 wherein the identified virtual address is currently used by an identified source-side port when part of the path is shared by two source nodes sending data to the same destination node.

104. The method of claim 92 including providing the identified virtual address to the source node for use in sending data to the destination node.

105. A network manager for establishing a path between a source node and a destination node through a network of switches, comprising:

- a component that identifies switches forming a path between the source node and the destination node;
- a component that identifies a virtual address for sending data from the source node to the destination node through the identified switches; and

a component that configures each of the identified switches to route data sent to the identified virtual address through the identified switches from the source node to the destination node.

106. The network manager of claim 105 including:

a component that identifies a virtual address for sending data from the destination node to the source node; and

a component that configures each of the identified switches to route data sent to the identified virtual address through the identified switches from the destination node to the source node.

107. The network manager of claim 105 including:

a component that identifies switches forming a path between the destination node and the source node.

108. The network manager of claim 107 wherein the path from the source node to the destination node includes one port that is not in the path from the destination node to the source node.

109. The network manager of claim 107 wherein the path from the source node to the destination node is different from the path from the destination node to the source node.

110. The network manager of claim 105 wherein each switch has ports with a mapping of virtual addresses to another port of the switch.

111. The network manager of claim 105 wherein when data is received at a port of a switch, the identified virtual address is used to retrieve an indication of another port of the switch through which the data is transmitted.

112. The network manager of claim 105 wherein a path is established between the source node and each of a plurality of destination nodes by identifying ports of switches for each path.

113. The network manager of claim 105 wherein the data is a Fibre Channel frame.

114. The network manager of claim 105 wherein the data is an InfiniBand frame.

115. The network manager of claim 105 wherein the switches are interconnect fabric modules.

116. The network manager of claim 105 wherein when a port of a switch receives data with a virtual address that has not been set for the port, the port does not forward the data.

117. The network manager of claim 105 wherein each switch has a source-side port and the identified virtual address is not currently used by any of the source-side ports.

118. The network manager of claim 105 wherein each switch has a source-side port and the identified virtual address is currently used by a source-side port when part of the path is shared by two source nodes sending data to the same destination node.

119. A network manager for establishing a path between a source node and a destination node through a network of routing devices, comprising:

means for identifying ports of routing devices forming a path between the source node and the destination node, each routing device of the path having an identified source-side port and an identified destination-side port;

means for identifying a virtual address for sending data from the source node to the destination node; and

setting each of the identified source-side ports to route data sent to the identified virtual address through the identified destination-side port of the routing device.

120. The network manager of claim 119 including:
means for identifying a virtual address for sending data from the destination node to the source node; and
means for setting each of the identified destination-side ports to route data sent to the identified virtual address through the identified source-side port of the routing device.

121. The network manager of claim 119 wherein a routing device is a switch.

122. The network manager of claim 119 wherein each port of each routing device has a means for mapping virtual addresses to another port of the routing device.

123. The network manager of claim 119 including means for, when data is received at a port of a routing device, retrieving an indication of another port using the identified virtual address and sending the data out of the routing device through the other port.

124. The network manager of claim 119 including means for establishing a path between the source node and each of a plurality of destination nodes by identifying ports of routing devices for each path.

125. The network manager of claim 119 wherein the data is a Fibre Channel frame.

126. The network manager of claim 119 wherein the data is an InfiniBand frame.

127. A method in a switch for transmitting frames to a network manager, the method comprising:

receiving a frame having a destination virtual address; and
upon receiving the frame,

determining whether the destination virtual address of the frame is reserved;
when the destination virtual address of the frame is reserved,
determining whether another virtual address of the frame maps to a port of the switch;
when the other virtual address of the frame maps to a port of the switch, transmitting the frame via the mapped-to port; and
when the other virtual address of the frame does not map to a port of the switch, transmitting the frame to the network manager.

128. The method of claim 127 wherein the destination virtual address and the other virtual address are stored in a header of the frame.

129. The method of claim 127 wherein the determining of whether another virtual address of the frame maps to a port of a switch includes checking a mapping of virtual addresses to ports.

130. The method of claim 129 wherein each port of the switch has its own mapping.

131. The method of claim 127 wherein the transmitting of the frame via the mapped-to port transmits the frame to the network manager.

132. The method of claim 127 wherein the mapped-to port transmits the frame via in-band communications.

133. The method of claim 127 wherein the network manager transmits the frame via out-of-band communications.

134. The method of claim 127 wherein the network manager is distributed to devices connected to switches and the network manager transmits the frame via an out-of-band communications to a device connected to the switch.

135. The method of claim 127 wherein the network manager is centralized and the frame is transmitted to the network manager via in-band communications.

136. A method in a routing device for transmitting frames to a network manager, the method comprising:

receiving a frame having a virtual address; and
determining whether the virtual address of the frame is reserved;
when the virtual address of the frame is reserved, providing the frame to the network manager;
when the virtual address of the frame is not reserved, transmitting the frame via a port of the routing device based on a mapping of virtual addresses to ports.

137. The method of claim 136 wherein the providing of the frame to the network manager includes:

determining whether another virtual address of the frame maps to a port of the routing device,
when the other virtual address of the frame maps to a port of the routing device,
transmitting the frame via the mapped-to port; and
when the other virtual address of the frame does not map to a port of the switch,
transmitting the frame directly to the network manager.

138. The method of claim 137 wherein the virtual address and the other virtual address are stored in a header of the frame.

139. The method of claim 137 wherein the determining of whether another virtual address of the frame maps to a port of the routing device includes checking a mapping of virtual addresses to ports.

140. The method of claim 139 wherein each port of the routing device has its own mapping.

141. The method of claim 136 the routing device is a switch.

142. The method of claim 136 wherein the providing of the frame to the network manager transmits the frame via out-of-band communications.

143. The method of claim 136 wherein the network manager is distributed to devices connected to routing devices and the providing of the frame to the network manager transmits the frame via out-of-band communications to a device connected to the routing device.

144. A routing device for transmitting data to a network manager, comprising:
a component that receives data having a virtual address; and
a component that, when the virtual address of the data is reserved, provides the data to the network manager and that, when the virtual address is not reserved, transmits the data via a port of the routing device based on a mapping of virtual addresses to ports.

145. The routing device of claim 144 wherein the providing of the data to the network manager includes:

a component that, when the other virtual address of the data maps to a port of the routing device, transmits the data via the mapped-to port and, when the other virtual address does not map to a port of the routing device, transmits the frame directly to the network manager.

146. The routing device of claim 145 wherein the virtual address and the other virtual address are stored in a header of the data.

147. The routing device of claim 146 wherein each port of the routing device has its own mapping.

148. The routing device of claim 144 the routing device is a switch.

149. The routing device of claim 144 wherein the providing of the data to the network manager transmits the data via out-of-band communications.

150. The routing device of claim 144 wherein the network manager is distributed to devices connected to routing devices and the providing of the data to the network manager transmits the data via an out-of-band communications to a device connected to the routing device.

151. A switch for transmitting data to a network manager, comprising:
means for receiving data having a virtual address; and
means for providing the data to the network manager when the virtual address of the data is reserved and for transmitting the data via a port of the routing device based on a mapping of virtual addresses to ports when the virtual address of the data is not reserved.

152. The switch of claim 151 wherein the providing of the data to the network manager includes:

means for transmitting the data via a mapped-to port when the virtual address of the data maps to a port of the switch and for transmitting the data directly to the network manager when the other virtual address of the data does not map to a port of the switch.

153. The switch of claim 151 wherein the virtual address and the other virtual address are stored in a header of the data.

154. The switch of claim 153 wherein each port of the switch has its own mapping.

155. The switch of claim 154 wherein the providing of the data to the network manager transmits the data via out-of-band communications.

156. The switch of claim 151 wherein the network manager is distributed to devices connected to switches and the providing of the data to the network manager transmits the data via out-of-band communication to a device connected to the switch.




Fig. 1





Fig. 2

*Fig. 3*




Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8